

On the Power of Many One-Bit Provers

Per Austrin*

Aalto University and KTH Royal Institute of Technology

Johan Håstad†

KTH Royal Institute of Technology

Rafael Pass‡

Cornell University

January 15, 2013

Abstract

We study the class of languages, denoted by $\text{MIP1}[k, 1 - \epsilon, s]$, which have k -prover games where each prover just sends a *single* bit, with completeness $1 - \epsilon$ and soundness error s . For the case that $k = 1$ (i.e., for the case of interactive proofs), Goldreich, Vadhan and Wigderson (*Computational Complexity'02*) demonstrate that SZK exactly characterizes languages having 1-bit proof systems with “non-trivial” soundness (i.e., $1/2 < s \leq 1 - 2\epsilon$). We demonstrate that for the case that $k \geq 2$, 1-bit k -prover games exhibit a significantly richer structure:

- (Folklore) When $s \leq \frac{1}{2^k} - \epsilon$, $\text{MIP1}[k, 1 - \epsilon, s] = \text{BPP}$;
- When $\frac{1}{2^k} + \epsilon \leq s < \frac{2}{2^k} - \epsilon$, $\text{MIP1}[k, 1 - \epsilon, s] = \text{SZK}$;
- When $s \geq \frac{2}{2^k} + \epsilon$, $\text{AM} \subseteq \text{MIP1}[k, 1 - \epsilon, s]$;
- For $s \leq 0.62k/2^k$ and sufficiently large k , $\text{MIP1}[k, 1 - \epsilon, s] \subseteq \text{EXP}$;
- For $s \geq 2k/2^k$, $\text{MIP1}[k, 1, 1 - \epsilon, s] = \text{NEXP}$.

As such, 1-bit k -prover games yield a natural “quantitative” approach to relating complexity classes such as $\text{BPP}, \text{SZK}, \text{AM}, \text{EXP}$, and NEXP . We leave open the question of whether a more fine-grained hierarchy (between AM and NEXP) can be established for the case when $s \geq \frac{2}{2^k} + \epsilon$.

1 Introduction

We study the expressiveness of k -prover games (introduced by Ben-Or, Goldwasser, Kilian and Wigderson [BOGKW88]), where each prover sends a *single* bit. Let $\text{MIP1}[k, 1 - \epsilon, s]$ denote the class of languages having a k -prover game where each prover sends a single bit, completeness $1 - \epsilon$, and soundness error s . Throughout the paper, we think of k as a constant and ϵ as an arbitrarily small constant. Clearly, for a fixed k , as s increases the corresponding complexity class can only become larger. We are interested in understanding to what extent the complexity class grows, and whether the growth is “smooth” or if threshold phenomena occur.

When the soundness error is “too small”, only trivial languages can have such games. In particular, provers sending random bits succeed with probability at least $(1 - \epsilon)2^{-k}$, placing the language of any protocol with smaller soundness in BPP .

*E-mail: austrin@kth.se.

†E-mail: johanh@kth.se. Funded by ERC advanced grant 226202.

‡E-mail: rafael@cs.cornell.edu. Funded by NSF, DARPA, AFOSR, a Sloan Foundation Fellowship, and a Microsoft Research Fellowship.

Theorem 1.1 (Folklore, implicit in [BGS98]). *For every $k \geq 1$, $\epsilon > 0$, we have*

$$\text{MIP1}[k, 1 - \epsilon, 1/2^k - \epsilon] = \text{BPP}$$

An interesting result by Goldreich, Vadhan and Wigderson [GVW02] shows that when $k = 1$ (i.e., for interactive proofs [GMR89, BM88]), whenever the soundness is “non-trivial”, then $\text{MIP1}[1, 1 - \epsilon, s]$ characterizes SZK , the class of languages having statistical zero-knowledge proofs. We here focus on the case when $k \geq 2$. As we shall see, in this setting, 1-bit k -prover games contains a richer variety of complexity classes. We take a first step towards characterizing these classes.

Our first result is a simple generalization of the result of [GVW02]: we show that when $\frac{1}{2^k} + \epsilon \leq s < \frac{2}{2^k} - \epsilon$, then $\text{MIP1}[k, 1 - \epsilon, s]$ characterizes SZK .

Theorem 1.2. *For every $k \geq 2$, $\epsilon > 0$, and $1/2^k + \epsilon < s < 2/2^k - \epsilon$, we have*

$$\text{MIP1}[k, 1 - \epsilon, s] = \text{SZK}.$$

Our main result next shows that when the soundness becomes just slightly higher than $2/2^k$, MIP1 s appear to become significantly more powerful; in particular, they contain all of AM .

Theorem 1.3 (Main Theorem). *For every $k \geq 2$ and $\epsilon > 0$*

$$\text{AM} \subseteq \text{MIP1}[k, 1 - \epsilon, 2/2^k + \epsilon]$$

For instance, when $k = 2$, our MIP1 has soundness error $\frac{1}{2} + \epsilon$. This result should be compared to Håstad’s 3-bit PCP [Hås01] that achieves the same soundness error. Since every 1-bit k -prover game yields a k -bit PCP, our MIP1 yields a 2-bit PCP for AM with soundness error $1/2 + \epsilon$; in contrast, the PCP resulting from our MIP1 is exponentially long, whereas Håstad’s PCP is polynomially long. Nonetheless, as we shall see shortly, our MIP1 construction heavily relies on Håstad’s PCP.

We leave open the question of whether $\text{MIP1}[k, 1 - \epsilon, 2/2^k + \epsilon]$ contains even richer complexity classes than AM . As a first step towards this question, we note that EXP is an upper bound on this class.

Theorem 1.4. *For all sufficiently large k , $\epsilon > 0$, $s \leq \frac{0.62k}{2^k}(1 - \epsilon)$ we have*

$$\text{MIP1}[k, 1 - \epsilon, s] \subseteq \text{EXP}.$$

This holds also for $k = 3$ and $s \leq 1/2 - \epsilon$.

Finally, we prove that for $k \geq 3$ and sufficiently high soundness error, k -prover 1-bit MIP1 s capture all of NEXP . This follows by using the PCP analogue of the classic $\text{MIP} = \text{NEXP}$ result [BFL91]. We sharpen the parameters by using more modern PCP machinery and then observing that the PCPs we use can be turned in to MIP1 at no cost. In particular using the recent results by Chan [Cha12], we get

Theorem 1.5. *For every $\epsilon > 0$ and $s = 2^{\lceil \log(k+1) \rceil}/2^k + \epsilon \leq 2k/2^k + \epsilon$,*

$$\text{MIP1}[k, 1 - \epsilon, s] = \text{NEXP}.$$

Taken together, these results demonstrate that k -prover games provide a natural “quantitative” way to relate complexity classes such as BPP , SZK , AM , EXP and NEXP . We leave open the question of whether $\text{MIP1}[k, 1 - \epsilon, s]$ contains an even more fine grained hierarchy of complexity classes in the regime where $s \geq 2/2^k + \epsilon$.

1.1 Related work

The work most closely related to ours is the work by Goldreich, Vadhan and Wigderson [GVW02] mentioned above which in turn builds on a work by Goldreich and Håstad [GH98]; just as we do, both these works investigate the complexity of interactive proofs with “laconic” provers. We have taken the question to an extreme in one direction (namely we focus only on provers that send a single bit); on the other hand, we have generalized the question by considering multi-prover interactive proofs, rather than just a single prover (as is the main focus in the above-mentioned works).

The large literature on PCP characterizations of NP (e.g., [AS98, ALM⁺98, BGLR94, BGS98, GLST98, ST00] and many others) is clearly also very related. As mentioned, a k -prover MIP1 yields a k -query PCP with the same soundness error, but of exponential length; typically, the PCP literature focuses on polynomial-length proofs. Nonetheless, we rely on both PCPs and techniques from this literature (most notably Fourier analysis) to analyze our proof system.

We also mention the recent work by Drucker [Dru11] that provides a PCP-type characterization of AM; his result is incomparable to our main theorem as he focuses on polynomial-length PCP proofs.

1.2 Outline

In Section 2 we present some definitions and background material that we use. In Section 3 we prove Theorem 1.2 for the SZK range. Our main result Theorem 1.3 is proved in Section 4. The Theorems 1.4 and 1.5 are proved in Section 5. Finally, we end with discussing some avenues for future work in Section 6.

2 Preliminaries

2.1 Laconic Proof systems

We assume familiarity with multi-prover interactive proofs and probabilistically checkable proofs.

Definition 2.1. $\text{IP}[k, c, s]$ denotes the class of problems having an two message protocol where the first message is sent by the Verifier and where the prover sends at most k bits and where the proof has soundness s and completeness c .

Definition 2.2. $\text{MIP1}[k, c, s]$ denotes the set of languages having a Multi-prover Interactive Proof System with k provers, each sending a single bit, soundness s , completeness c . The questions to the k provers are asked simultaneously. In other words, all questions are formulated before any answer is received.

Fact 2.3. For every $k \geq 1$, $0 \leq s < c \leq 1$, we have

$$\text{IP}[k, c, s] \subseteq \text{MIP1}[k, c, s].$$

When constructing MIP1 it is convenient to rely on efficient PCPs. There are general translations from PCPs to MIP1s (one is given in [BGS98]) if one accepts a slight loss in the parameters. In the cases we are interested in, however, by a slight extension of the analysis we can turn the PCP directly into a MIP1 without any loss in parameters.

2.2 Statistical Zero Knowledge

For our characterization of the SZK range, we only need to rely on the following result of [GVW02] relating SZK to laconic IP systems.

Theorem 2.4 ([GVW02], Theorem 3.1). *For every c, s such that $1 > c^2 > s > c/2 > 0$, it holds that $\text{IP}[1, c, s] = \text{SZK}$.*

2.3 Fourier Analysis of Boolean Functions

For two vectors $x, y \in \{0, 1\}^n$ we write $x \oplus y$ for their pointwise sum modulo 2. Given $a \in \{0, 1\}^n$ we write $\chi_a : \{0, 1\}^n \rightarrow \mathbb{R}$ for the *character* (which is in fact a linear function) $\chi_a(x) = (-1)^{\sum_{i=1}^n a_i x_i}$.

Any Boolean function $f : \{0, 1\}^n \rightarrow \mathbb{R}$ can be uniquely decomposed as a linear combination of characters

$$f(x) = \sum_{a \in \{0, 1\}^n} \hat{f}(a) \chi_a(x),$$

where $\hat{f}(a) = \mathbb{E}_x[f(x)\chi_a(x)]$ are the *Fourier coefficients* of f .

We recall Plancherel's equality: for any $f : \{0, 1\}^n \rightarrow \mathbb{R}$, we have

$$\sum_a \hat{f}(a)^2 = \mathbb{E}_x[f(x)^2].$$

2.4 Inapproximability of Linear Equations

Our proof system for AM is based on the optimal inapproximability result for linear equations mod 2 by Håstad [Hås01], defined next.

Definition 2.5. An instance Ψ of MAX 3-LIN-2 consists of a set of equations in n variables x_1, \dots, x_n over $\{0, 1\}$. Each equation is of the form $\chi_l(x) = b$ for some $l \in \{0, 1\}^n$ of weight 3 and some $b \in \{-1, 1\}$. We denote by $\text{Opt}(\Psi) \in [0, 1]$ the maximum fraction of equations satisfied by any assignment to x .

Theorem 2.6 ([Hås01]). *For every $\epsilon > 0$, given a MAX 3-LIN-2 instance Ψ , it is NP-hard to determine whether $\text{Opt}(\Psi) \leq 1 - \epsilon$ or whether $\text{Opt}(\Psi) \geq \frac{1+\epsilon}{2}$.*

3 The SZK range

Theorem 3.1. *For every $k \geq 1, \epsilon > 0$, we have*

$$\text{IP}[k, 1 - \epsilon, 1/2^k + \epsilon] \supseteq \text{SZK}.$$

Proof. Follows by repetition of the protocol from Theorem 2.4 and the fact that there is no problem with parallel repetition for one-prover proof systems. \square

Proposition 3.2. *For every $k \geq 1, 0 \leq s \leq c \leq 1$, we have*

$$\text{MIP1}[k, c, s] \subseteq \text{IP}[1, c, 2^{k-1}s].$$

Proof. Given a MIP1 protocol (V, P_1, \dots, P_k) for a language L , we construct a single-prover protocol (V', P') as follows. The verifier V' runs V to generate k messages x_1, \dots, x_k , and sends x_1 to the prover P' . The prover P' acts as P_1 and responds with an answer $y_1 \in \{0, 1\}$. V' accepts iff there are bits y_2, \dots, y_k such that the original verifier V accepts on the transcript $(x_1, \dots, x_k, y_1, \dots, y_k)$. Clearly, the completeness of (V', P') is at least that of the original protocol. For the soundness, suppose that there is a strategy for P' that makes the verifier accept with probability s' . Construct a strategy for the original protocol by having P_1 act as P' and P_2, \dots, P_k return random answers. Clearly, these provers make V accept with probability at least $s'/2^{k-1}$, implying $s' \leq 2^{k-1}s$ as desired. \square

Theorem 3.3. *For every $k \geq 1$, and every $\epsilon > 0$ it holds that*

$$\text{MIP1}[k, 1 - \epsilon, 2/2^k(1 - 2\epsilon)] \subseteq \text{SZK}$$

Proof. We have

$$\begin{aligned} \text{MIP1}[k, 1 - \epsilon, 2/2^k(1 - 2\epsilon)] &\subseteq \text{IP}[1, 1, 1 - \epsilon, 1 - 2\epsilon] && \text{(Proposition 3.2)} \\ &\subseteq \text{SZK} && \text{(Theorem 2.4)} \end{aligned}$$

□

4 Proof systems for AM

First we note that, at a cost of an arbitrarily small loss in soundness and completeness, we may restrict ourselves to proof systems for NP.

Lemma 4.1. *If $\text{NP} \subseteq \text{MIP1}[k, c, s]$ then for every $\epsilon > 0$ it holds that $\text{AM} \subseteq \text{MIP1}[k, c - \epsilon, s + \epsilon]$*

Proof. Let $L \in \text{AM}$. We remind the reader that this is equivalent to the existence of a language $L' \in \text{NP}$ such that $x \in L$ iff $(x, r) \in L'$ with high probability for a random string r (of an appropriate polynomial length). Without loss of generality, we may assume that the protocol for L has completeness $1 - \epsilon$ and soundness ϵ . The MIP1 verifier for L simply sends Arthur's random string r to each of the k provers and then executes the MIP1 protocol assumed to exist for $L' \in \text{NP}$.

If $x \in L$ then with probability $1 - \epsilon$ over r we have $(x, r) \in L'$ in which case the provers convince the verifier with probability $\geq c$.

On the other hand $x \notin L$ then the probability that the provers accept is at most $\Pr_r[(x, r) \in L'] + \Pr[(x, r) \notin L'] \Pr[\text{accept} | (x, r) \notin L'] \leq \epsilon + s$. □

4.1 Warm-up: the case of 2 provers

We start off with the case of only 2 provers, as this case is somewhat simpler than the general case, and will be used to obtain the general case.

Theorem 4.2. *For every $\epsilon > 0$*

$$\text{NP} \subseteq \text{MIP1}[2, 1 - \epsilon, 1/2 + \epsilon].$$

Proof. We reduce from the MAX 3-LIN-2 problem. Given is a MAX 3-LIN-2 instance Ψ , on n variables x_1, \dots, x_n and m linear equations $\{l_i(x_i) = b_i\}_{i \in [m]}$.

The provers are expected to provide oracle access to the Hadamard encoding of a $(1 - \epsilon)$ -satisfying assignment $x \in \{0, 1\}^n$. In other words, the verifier will give each prover a vector $a \in \{0, 1\}^n$ and expects in response the value of the linear function $\chi_a(x) \in \{-1, 1\}$.

The verifier proceeds as follows:

1. Pick a random equation $\chi_l(x) = b$ in Ψ
2. Pick random $y \in \{0, 1\}^n$
3. Check that $P_2(y) \cdot P_1(l \oplus y) = b$

It is easy to see that there is a strategy for the provers which makes the verifier accept with probability at least $\text{Opt}(\Psi)$. More interestingly, we will now prove that, $\text{Opt}(\Psi)$ is *exactly* the maximum acceptance probability, over any strategy for P_1 and P_2 .

We can then write the acceptance probability of the verifier as

$$\Pr[\text{Verifier accepts}] = \mathbb{E}_{\substack{y \in \{0,1\}^n \\ (l,b) \in \Psi}} \left[\frac{1 + bP_1(l \oplus y)P_2(y)}{2} \right]. \quad (1)$$

Replacing the two functions by their Fourier expansion we need to analyze

$$\sum_{a,a'} \hat{P}_1(a)\hat{P}_2(a') \mathbb{E}_{y,(l,b)} [b\chi_a(l \oplus y)\chi_{a'}(y)].$$

All terms with $a \neq a'$ have expectation 0 and furthermore we have

$$\left| \mathbb{E}_{(l,b)} [b\chi_a(l)] \right| \leq 2 \text{Opt}(\Psi) - 1,$$

as the assignment given by a satisfies at most an $\text{Opt}(\Psi)$ fraction of the equations and at least a fraction $1 - \text{Opt}(\Psi)$ as its negation does not satisfy more than a $\text{Opt}(\Psi)$ fraction. We conclude that (1) is bounded by

$$\frac{1 + \sum_a |\hat{P}_1(a)\hat{P}_2(a)| (2 \text{Opt}(\Psi) - 1)}{2}.$$

Finally note that, by Cauchy-Schwarz,

$$\sum_a |\hat{P}_1(a)\hat{P}_2(a)| \leq \left(\sum_a \hat{P}_1^2(a) \right)^{1/2} \left(\sum_a \hat{P}_2^2(a) \right)^{1/2} = 1$$

and this finishes the argument. \square

4.2 The general case

We have

Theorem 4.3. *For every $k \geq 2$, $\epsilon > 0$, we have*

$$\text{NP} \subseteq \text{MIP1}[k, 1 - \epsilon, 2/2^k + \epsilon].$$

Proof. As before, we design a MIP1 system for linear equations. Given is a MAX 3-LIN-2 instance Ψ , in which either $\text{Opt}(\Psi) \geq 1 - \epsilon_0$, or $\text{Opt}(\Psi) \leq \frac{1+\epsilon_0}{2}$, where ϵ_0 will be chosen small enough to get the completeness and soundness bound that we want.

The verifier again expects all the k provers to provide answers to the Hadamard coding of the good assignment, and it then does the obvious generalization of the $k = 2$ case:

1. Pick $k - 1$ random equations $l_j(x) = b_j$, $1 \leq j \leq k - 1$
2. Pick random $y \in \{0,1\}^n$
3. Check that $P_j(l_j \oplus y) \cdot P_k(y) = b_j$ for every $1 \leq j \leq k - 1$

It is clear that the completeness is at least $(1 - \epsilon_0)^{k-1} \geq 1 - k\epsilon_0$. Thus, as long as $\epsilon_0 \leq \epsilon/k$, we have the desired completeness.

Let us now study the soundness, i.e., the maximum possible acceptance probability of verifier, given that $\text{Opt}(\Psi) \leq \frac{1+\epsilon_0}{2}$.

We say that prover P_j succeeds if $P_j(l_j \oplus y) \cdot P_k(y) = b_j$. From the analysis of the previous theorem, we know that the probability that P_j succeeds is at most $\frac{1+\epsilon_0}{2}$. Thus, if the events that the different provers succeed were independent, we would obtain the desired soundness of $\approx 2^{1-k}$. However, *a priori*, it may be that the success events of the provers are very correlated, e.g., it could be that if one succeeds then they all succeed.

To cope with this, we need to obtain a more robust version of the previous analysis. Let $\frac{1+\delta_j(y)}{2}$ be the probability that P_j succeeds given that y is chosen. We have the following lemma.

Lemma 4.4. $\mathbb{E}_y[\delta_j^2(y)] \leq \epsilon_0^2$.

Proof. We have $\delta_j(y) = \mathbb{E}_{(l,b)}[bP_k(y)P_j(l \oplus y)]$ and thus

$$\mathbb{E}_y[\delta_j^2(y)] = \mathbb{E}_{(l,b),(l',b'),y}[bb'P_j(l \oplus y)P_j(l' \oplus y)].$$

Similarly to the case $k = 2$ we replace the function by its Fourier expansion and we are left to analyze

$$\sum_{a,a'} \hat{P}_j(a)\hat{P}_j(a') \mathbb{E}_{y,(l,b),(l',b')}[bb'\chi_a(l+y)\chi_{a'}(l'+y)].$$

Again we only have nonzero terms when $a = a'$. For these terms it is easy to see that

$$\left| \mathbb{E}_{(l,b),(l',b')}[bb'\chi_a(l)\chi_a(l')] \right| \leq (2\text{Opt}(\Psi) - 1)^2 = \epsilon_0^2.$$

Using $\sum_a \hat{P}_j(a)^2 = 1$, the lemma follows. \square

Lemma 4.4 implies that the fraction of y such that $\delta_j(y) \geq \sqrt{\epsilon_0}$ is bounded by ϵ_0 .

We conclude that the, for the y chosen, the probability that $\delta_j(y) \geq \sqrt{\epsilon_0}$ for any j is bounded by $k\epsilon_0$. On the other hand if $\delta_j(y) \leq \sqrt{\epsilon_0}$ for all values of j the probability that the verifier accepts is bounded by $(\frac{1+\sqrt{\epsilon_0}}{2})^{k-1}$. We conclude that the overall probability that the verifier accepts is bounded by

$$k\epsilon_0 + \left(\frac{1 + \sqrt{\epsilon_0}}{2} \right)^{k-1},$$

and choosing ϵ_0 sufficiently small, this is bounded by $2^{1-k} + \epsilon$. \square

5 The High End – EXP and NEXP Results

In this section we prove Theorems 1.4 and 1.5. These are essentially just “blow-ups” of corresponding approximation algorithms and inapproximability results.

Theorem 5.1 (Theorem 1.4 restated). *For all sufficiently large k , $\epsilon > 0$, $s \leq \frac{0.62k}{2^k}(1-\epsilon)$ we have*

$$\text{MIP1}[k, 1 - \epsilon, s] \subseteq \text{EXP}.$$

This holds also for $k = 3$ and $s \leq 1/2 - \epsilon$.

Proof. Let $L \in \text{MIP1}[k, 1 - \epsilon, s]$ with $s \leq \frac{0.62k}{2^k}(1 - \epsilon)$. Given an instance, the task of determining whether $x \in L$ boils down to finding the best joint strategy for the k provers. If the verifier uses r random bits she can send at most 2^r different queries to each prover, thus the optimal strategy can be described by $k \cdot 2^r = 2^{\text{poly}(|x|)}$ bits. Further, for each outcome of the verifier's randomness, the acceptance criterion is a constraint on some k bits of the strategy. In other words, what we have is an exponentially large Max k -CSP instance. The value of this instance can be approximated in time polynomial in its size to within a factor $0.62k/2^k$ by the algorithm of Makarychev and Makarychev [MM12]. For the case $k = 3$ we use the $1/2$ -approximate Max 3-CSP algorithm of Zwick [Zwi98]. \square

Next we show that if the soundness is sufficiently large, exponential-size k -query PCP systems can express every language in **NEXP**.

Theorem 5.2. *For $t = 2^{\lceil \log_2(k+1) \rceil}$ ($k+1$ rounded up to the next power of 2) we have*

$$\text{MIP1}[k, 1 - \epsilon, t/2^k + \epsilon] = \text{NEXP}.$$

This immediately implies Theorem 1.5.

Proof sketch. The proof follows from a upscaling of the recent PCP of Chan [Cha12] that gives a predicate of arity k which has t accepting configurations and which is approximation resistant.

In a standard PCP, the verifier runs in polynomial time, uses a logarithmic number of random coins and reads a constant number of bits in a polynomial size proof and verifies an NP-statement. We are currently interested in the situation where the crucial parameters, except the running time of the verifier, are exponentially larger.

To be more precise we are interested in a polynomial time verifier, that uses a polynomial number of random coins and gets one bit each from k different provers that respond to questions of polynomial length.

As is convenient for us, Chan already analyzed his PCP in the k -partite situation where each bit is read from a separate table. This model is exactly the same as a k -prover model and hence this difference is only syntactical.

It remains to address the question on how to make the upscaled verifier run in polynomial time. This amounts to saying that a verifier of an **NEXP** statement runs in polynomial time. This was explicitly needed in [BFL91] but this paper predates the PCP-Theorem. The fact that this is true also for upscaled versions of the PCP-Theorem has been explicitly stated in [BGS98] and [BSGH⁺05]. The intuitive reason that this is true is that the verifier only needs to ensure that some bits in a suitable encoding of the inputs are correct and this takes polynomial time in the size of the input but not the other parameters of the proof. \square

6 Concluding Remarks

There are a number of interesting avenues for further work. In this paper we focused solely on the case of almost perfect completeness and each prover sending exactly 1 bit. Obviously, understanding what happens with the expressiveness of these systems for other completeness values (in particular perfect completeness) and slightly less laconic provers would be very interesting. By simple extensions of the methods used in this paper it is possible to get some results but it would be interesting to see if perfect completeness could lead to a significantly different situation in any range of parameters.

There is also a specific question more directly related to the current paper. There is a huge gap between our lower bound **AM** and upper bound **EXP** for soundness $s = 2/2^k + \epsilon$. It seems quite plausible that an upper bound for this range of s should be **PSPACE** rather

than EXP – proving this essentially boils down to proving that there is a $\delta > 0$ such that bipartite instances of Max 2-CSP can be approximated within a factor $1/2 + \delta$ in polylog-space (and not necessarily polynomial time). We hope that the recent algorithms for Max Cut, in particular [KS11], can be adapted to achieve this.

Even if this turns out to be true, whether the correct class here is AM or PSPACE or something in between we have little intuition about.

Acknowledgment. We are grateful to Salil Vadhan for pointing out a simple proof of the lower bound given in Theorem 1.2 rather than the more complicated proof with worse parameters that we originally had. We are also grateful to Madhu Sudan and Or Meir for discussions on how to blow-up PCPs.

References

- [ALM⁺98] Sanjeev Arora, Carsten Lund, Rajeev Motwani, Madhu Sudan, and Mario Szegedy. Proof verification and the hardness of approximation problems. *J. ACM*, 45(3):501–555, 1998.
- [AM09] P. Austrin and E. Mossel. Approximation resistant predicates from pairwise independence. *Computational Complexity*, 18:249–271, 2009.
- [AS98] Sanjeev Arora and Shmuel Safra. Probabilistic checking of proofs: A new characterization of NP. *J. ACM*, 45(1):70–122, 1998.
- [BFL91] László Babai, Lance Fortnow, and Carsten Lund. Non-deterministic exponential time has two-prover interactive protocols. *Computational Complexity*, 1:3–40, 1991.
- [BGLR94] Mihir Bellare, Shafi Goldwasser, Carsten Lund, and Alexander Russell. Efficient probabilistic checkable proofs and applications to approximation. In *STOC*, page 820, 1994.
- [BGS98] Mihir Bellare, Oded Goldreich, and Madhu Sudan. Free bits, PCPs, and nonapproximability-towards tight results. *SIAM J. Comput.*, 27(3):804–915, 1998.
- [BM88] László Babai and Shlomo Moran. Arthur-Merlin games: A randomized proof system, and a hierarchy of complexity classes. *J. Comput. Syst. Sci.*, 36(2):254–276, 1988.
- [BOGKW88] Michael Ben-Or, Shafi Goldwasser, Joe Kilian, and Avi Wigderson. Multi-prover interactive proofs: How to remove intractability assumptions. In *STOC*, pages 113–131, 1988.
- [BSGH⁺05] Eli Ben-Sasson, Oded Goldreich, Prahladh Harsha, Madhu Sudan, and Salil Vadhan. Short PCPs verifiable in polylogarithmic time. In *Proceedings of the 20th Annual IEEE Conference on Computational Complexity*, CCC ’05, pages 120–134, Washington, DC, USA, 2005. IEEE Computer Society.
- [Cha12] Siu On Chan. Approximation Resistance from Pairwise Independent Subgroups. *Electronic Colloquium on Computational Complexity (ECCC)*, 19:110, 2012.
- [Dru11] Andrew Drucker. A PCP characterization of AM. In *ICALP (1)*, pages 581–592, 2011.

[GH98] Oded Goldreich and Johan Håstad. On the complexity of interactive proofs with bounded communication. *Inf. Process. Lett.*, 67(4):205–214, 1998.

[GLST98] Venkatesan Guruswami, Daniel Lewin, Madhu Sudan, and Luca Trevisan. A tight characterization of np with 3 query PCPs. In *FOCS*, pages 8–17, 1998.

[GMR89] Shafi Goldwasser, Silvio Micali, and Charles Rackoff. The knowledge complexity of interactive proof systems. *SIAM J. Comput.*, 18(1):186–208, 1989.

[GVW02] Oded Goldreich, Salil P. Vadhan, and Avi Wigderson. On interactive proofs with a laconic prover. *Computational Complexity*, 11(1-2):1–53, 2002.

[Hås01] Johan Håstad. Some Optimal Inapproximability Results. *Journal of the ACM*, 48(4):798–859, 2001.

[Kho02] S. Khot. On the power of unique 2-prover 1-round games. In *Proceedings of 34th ACM Symposium on Theory of Computing*, pages 767–775, 2002.

[KS11] Satyen Kale and C. Seshadhri. Combinatorial approximation algorithms for maxcut using random walks. In *ICS*, pages 367–388, 2011.

[MM12] Konstantin Makarychev and Yury Makarychev. Approximation Algorithm for Non-Boolean MAX k-CSP. To appear in APPROX, 2012.

[ST00] Alex Samorodnitsky and Luca Trevisan. A PCP characterization of NP with optimal amortized query complexity. In *STOC*, pages 191–199, 2000.

[Zwi98] Uri Zwick. Approximation algorithms for constraint satisfaction problems involving at most three variables per constraint. In *SODA*, pages 201–210, 1998.